



Computation of Optimal Transport and Related Hedging Problems via Penalization and Neural Networks

Stephan Eckstein¹ · Michael Kupper¹

Published online: 20 February 2019
© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

This paper presents a widely applicable approach to solving (multi-marginal, martingale) optimal transport and related problems via neural networks. The core idea is to penalize the optimization problem in its dual formulation and reduce it to a finite dimensional one which corresponds to optimizing a neural network with smooth objective function. We present numerical examples from optimal transport, martingale optimal transport, portfolio optimization under uncertainty and generative adversarial networks that showcase the generality and effectiveness of the approach.

Keywords Optimal transport · Robust hedging · Numerical method · Duality · Regularisation · Feedforward networks · Knightian uncertainty · Distributional robustness

1 Introduction

In this paper we present a penalization method which allows to compute a wide class of optimization problems of the form

$$\phi(f) = \sup_{\nu \in \mathcal{Q}} \int f d\nu$$

by means of neural networks. The most widely known representative of such a functional occurs in the optimal transport problem, to be introduced shortly. More generally,

We thank Daniel Bartl, Fabian Both, Jens Jackwerth, Mathias Pohl, Stefan Volkwein and two anonymous referees for helpful comments.

✉ Michael Kupper
kupper@uni-konstanz.de

Stephan Eckstein
stephan.eckstein@uni-konstanz.de

¹ Department Mathematics and Statistics, University of Konstanz, Constance, Germany

these functionals appear for instance in the representation of coherent risk measures [4] as the worst-case expected loss over a class \mathcal{Q} of scenario probabilities, in the representation of nonlinear expectations [40], or as the upper bound of arbitrage-free prices for a contingent claim f , see e.g. [24]. To solve the initial problem $\phi(f)$ we will make use of its dual formulation and restrict to the subclass of those optimization problems which can be realized as a minimal superhedging price

$$\phi(f) = \inf_{\substack{h \in \mathcal{H}: \\ h \geq f}} \int h d\mu_0.$$

for some $\mu_0 \in \mathcal{Q}$, where \mathcal{H} is a set of continuous and bounded functions $h : \mathcal{X} \rightarrow \mathbb{R}$, where the relation of \mathcal{H} and \mathcal{Q} is given at the beginning of Sect. 2. A very similar class of optimization problems in an abstract framework of Banach lattices is studied in [20]. Under sufficient regularity conditions the values of the primal problem $\sup_{\nu \in \mathcal{Q}} \int f d\nu$ and its dual problem $\inf_{h \in \mathcal{H}: h \geq f} \int h d\mu_0$ can be shown to coincide, see e.g. [16] for related pricing-hedging dualities.

A typical example is the Kantorovich relaxation [35] of Monge's optimal transport problem, where \mathcal{Q} is the set of probability measures on a product space $\mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2$ with given marginals μ_1 and μ_2 , and where \mathcal{H} is the set of all continuous and bounded functions $h(x_1, x_2) = h_1(x_1) + h_2(x_2)$ and $\int h d\mu_0 = \int_{\mathcal{X}_1} h_1 d\mu_1 + \int_{\mathcal{X}_2} h_2 d\mu_2$. Further frequently studied problems in this class include multi-marginal optimal transport and Wasserstein distances (see e.g. [5,49,50]), martingale optimal transport (see e.g. [7,25,30,32]), value at risk under dependence uncertainty (see e.g. [11,21,43]), or calculating worst case copula values and improved Fréchet–Hoeffding bounds (see e.g. [6,39]). Moreover, $\phi(f)$ serves as a building block for several other problems, like generative adversarial networks (where additionally, the optimization includes generating a distribution, see e.g. [3,22,29]), portfolio choice under dependence uncertainty (where additionally, portfolio weights are optimized, see e.g. [10,42]), or robust optimized certainty equivalents (see e.g. [19]). In these cases, the solution approach presented in this paper is still applicable.

Summary of the approach The goal is to solve $\phi(f)$ numerically. The implementation will build on the dual representation of $\phi(f)$. The first step is to go over to a finite dimensional setting, where the set \mathcal{H} is replaced by a subset \mathcal{H}^m :

$$\phi^m(f) = \inf_{\substack{h \in \mathcal{H}^m: \\ h \geq f}} \int h d\mu_0$$

Theoretically, we will look at a sequence $(\mathcal{H}^m)_{m \in \mathbb{N}}$ with $\mathcal{H}^1 \subseteq \mathcal{H}^2 \subseteq \dots \subseteq \mathcal{H}$ such that $\mathcal{H}^\infty := \cup_{m \in \mathbb{N}} \mathcal{H}^m$ is in a certain sense dense in \mathcal{H} . More concretely, \mathcal{H}^m can be a set of neural networks with a fixed structure (but unspecified parameter values), and m measures the number of neurons per layer.

To allow for a step-wise updating of the parameters (e.g. by gradient descent methods) for the space \mathcal{H}^m , the inequality constraint $h \geq f$ is penalized. To this end, we introduce a *reference probability measure* θ on the state space \mathcal{X} . Intuitively, this




Fig. 1 Occurring problems and their relations. The depicted convergences are studied in Sect. 2 and, in a more specific context of neural networks, in Sect. 3

measure will be used to sample points at which the inequality constraint $h \geq f$ can be tested. Further, we introduce a differentiable and nondecreasing *penalty function* $\beta : \mathbb{R} \rightarrow \mathbb{R}_+$. This leads to the penalized problem

$$\phi_{\theta,\beta}^m(f) = \inf_{h \in \mathcal{H}^m} \left\{ \int h \, d\mu_0 + \int \beta(f - h) \, d\theta \right\}.$$

For theoretical considerations we also introduce

$$\phi_{\theta,\beta}(f) = \inf_{h \in \mathcal{H}} \left\{ \int h \, d\mu_0 + \int \beta(f - h) \, d\theta \right\}.$$

We will again consider sequences of penalty functions $(\beta_\gamma)_{\gamma>0}$ parametrized by a penalty factor γ , and use the notation $\phi_{\theta,\gamma}(f) := \phi_{\theta,\beta_\gamma}(f)$ and $\phi_{\theta,\gamma}^m(f) := \phi_{\theta,\beta_\gamma}^m(f)$. Here, an increasing penalty factor can be seen as a more and more precise enforcing of the inequality constraint $h \geq f$.

The problems $\phi_{\theta,\gamma}^m(f)$ are the ones which are solved numerically. Chapters 2 and 3 study the relation between this problem which is eventually implemented, and the initial problem $\phi(f)$. To this end, we analyse how the introduced approximative problems behave for $m \rightarrow \infty$ and $\gamma \rightarrow \infty$. Figure 1 summarizes the occurring problems and their relations. Notably, we are only interested in convergence of optimal values, not that of optimizers.

The final step is to find a numerical solution of $\phi_{\theta,\gamma}^m(f)$, which means in practice finding the optimal parameters of the network \mathcal{H}^m . We use Tensorflow [1] and the Adam optimizer [37] to this end, and thus mostly regard this step as a black box. We will denote the numerical optimal solution by $\hat{\phi}_{\theta,\gamma}^m(f)$.

Implementation method: Related literature Penalization of optimal transport problems has been studied in several works (see e.g. [9,14,17,18,26,29,44,45,47]). Entropic penalization in particular is applied often, which is in close relation to the Schrödinger problem [38]. Cominetti and San Martín’s work [17] from 1994 on entropic penalization of arbitrary linear programs can be applied to purely discrete optimal transport. The basic idea in [17] is to obtain a strictly convex problem through penalization

which can be solved quicker and converges to the initial problem, for an increasing penalty factor. More recently, Cuturi [18] gives an efficient algorithm to compute discrete optimal transport problems with two marginals based on entropic penalization and Sinkhorn's matrix scaling algorithm. Genevay et al. [26] and Solomon et al. [47] go further in this direction and give algorithms to compute arbitrary optimal transport problems with two marginals, where the algorithm (for the case of continuous marginals) is based on a reproducing kernel Hilbert space approach, and discretization, respectively. In [26] the authors already mention that more general regularizations beyond the entropic one are possible. Among others Benamou et al. [9] and Schmitzer [44] use scaling algorithms related to [18] for a larger class of problems, including for example (discrete) multi-marginal, constrained and unbalanced optimal transport. Carlier et al. [14] show Γ -convergence of the entropic penalized Wasserstein-2 distance to the unpenalized one. The same kind of Γ -convergence is also subject of the studies related to the Schrödinger problem [38], even for more general cost functions. Recent research by Arjovsky et al. [3,29] inspired by generative adversarial networks include solving a particular optimal transport problem (the Wasserstein-1 distance) based on L^2 penalization. In these works, the numerical approach to solve optimal transport problems by parametrization of the dual variables by neural networks originated. Seguy et al. [45] apply a neural network based approach to arbitrary optimal transport problems with two marginals. Their theoretical results are broadly based on entropic penalization, discretization, and weakly continuous dependence of the optimal transport problem on the marginals.

Contribution The current paper gives a unifying numerical solution approach to problems of the form $\phi(f)$ based on penalization and neural networks. The focus lies both on general applicability with respect to the choice of problem, and also on a flexible framework regarding the solution method.

Compared to the existing literature, which often focusses on a single representative (often the optimal transport problem) among problems of the form $\phi(f)$, our theoretical results are widely applicable. Similarly, the penalization method and the resulting dual relations in this paper allow for many different forms of reference measure θ and penalty function β_γ , while the existing literature is often restricted to uniform or product reference measures, and exponential penalty function.¹ We show the effects of different reference measures and different penalty functions both theoretically in Theorem 2.2 and practically in the numerical examples in Sect. 4. In some examples the choice of an appropriate reference measure is crucial, see e.g. Sect. 4.4. Equation (2.6) of Theorem 2.2 also motivates an updating procedure for the reference measure to reduce the error arising from penalization, which is applied in Sect. 4.5.

The presented approach is showcased with several examples, which are mostly toy problems taken from existing papers. The reason we use toy problems is to allow for an evaluation of the numerical methods that can be based on analytical solutions.

¹ In discrete settings, the reference measure is usually the uniform distribution (see e.g. [18], where the penalization does not explicitly include a reference measure. The penalization applied is simply the entropy of a measure, which corresponds to the relative entropy with uniform reference measure). In non-discrete settings, usually the product measure of the marginals specified by the optimal transport problems are used (see e.g. [26,45]).

Structure of the paper In Sect. 2 we present the theoretical results on approximation and regularization. Section 3 discusses the particular case of \mathcal{H}^m as built by multi-layer feedforward networks. In Sect. 4 we illustrate the proposed method with several examples. All proofs are postponed to Sect. 5.

2 Regularization and Approximation of Hedging Functionals

Let $\mathcal{P}(\mathcal{X})$ be the set of all Borel probability measures on a Polish space \mathcal{X} , and denote by $C_b(\mathcal{X})$ the linear space of all continuous bounded functions $f : \mathcal{X} \rightarrow \mathbb{R}$. We consider the *superhedging functional*

$$\phi(f) := \inf \left\{ \int h \, d\mu_0 : h \geq f \text{ for some } h \in \mathcal{H} \right\} \tag{2.1}$$

for $f \in C_b(\mathcal{X})$, where $\mu_0 \in \mathcal{P}(\mathcal{X})$ is a pricing measure and $\mathcal{H} \subseteq C_b(\mathcal{X})$. Throughout this section we assume that \mathcal{H} is a linear space which contains the constants (i.e. the constant functions).

In order to derive a dual representation, we assume that ϕ is continuous from above, i.e. $\phi(f_n) \downarrow 0$ for every sequence (f_n) in $C_b(\mathcal{X})$ such that $f^n \downarrow 0$. By the nonlinear Daniell-Stone theorem it has a representation

$$\phi(f) = \max_{\mu \in \mathcal{Q}} \int f \, d\mu \tag{2.2}$$

for all $f \in C_b(\mathcal{X})$, and the nonempty set $\mathcal{Q} = \{ \mu \in \mathcal{P}(\mathcal{X}) : \int h \, d\mu = \int h \, d\mu_0 \text{ for all } h \in \mathcal{H} \}$. In particular $\mu_0 \in \mathcal{Q}$. The problems (2.2) and (2.1) are in duality and we refer to (2.2) as the primal and (2.1) as the dual formulation. For the details we refer to the Appendix A. There it is outlined how the duality extends to unbounded functions. However, for the sake of readability we focus on $C_b(\mathcal{X})$.

The following example illustrates the basic setting:

Example 2.1 Let $\mathcal{X} = \mathbb{R}^d$, and denote by $\Pi(\mu_1, \dots, \mu_d)$ the set of all $\mu \in \mathcal{P}(\mathbb{R}^d)$ with first marginal μ_1 , second marginal μ_2 , etc. In the following examples, under the assumption that $\mathcal{Q} \neq \emptyset$ it is straightforward to verify that the corresponding superhedging functional is continuous from above.

(a) (Multi-marginal) optimal transport [35,50]:

$$\begin{aligned} \mathcal{Q} &= \Pi(\mu_1, \dots, \mu_d), \\ \mathcal{H} &= \left\{ h \in C_b(\mathbb{R}^d) : h(x_1, \dots, x_d) = h_1(x_1) + \dots + h_d(x_d) \text{ for all } (x_1, \dots, x_d) \right. \\ &\quad \left. \in \mathbb{R}^d \text{ and some } h_i \in C_b(\mathbb{R}) \right\} \end{aligned}$$

(b) Martingale optimal transport [7,25]:

$$\begin{aligned} \mathcal{Q} &= \{ \mu \in \Pi(\mu_1, \dots, \mu_d) : \text{the canonical process on } \mathbb{R}^d \text{ is a } \mu\text{-martingale} \} \\ \mathcal{H} &= \left\{ h \in C_\kappa(\mathbb{R}^d) : h(x_1, \dots, x_d) = \sum_{i=1}^d h_i(x_i) \right. \\ &\quad \left. + \sum_{i=2}^d g_i(x_1, \dots, x_{i-1}) \cdot (x_i - x_{i-1}) \text{ for all } (x_1, \dots, x_d) \right. \\ &\quad \left. \in \mathbb{R}^d \text{ and some } h_i \in C_b(\mathbb{R}) \text{ and } g_i \in C_b(\mathbb{R}^{i-1}) \right\} \end{aligned}$$

where $C_\kappa(\mathbb{R}^d)$ denotes the space of all continuous functions of linear growth corresponding to $\kappa(x) := 1 + |x|$, see Appendix A. By Strassen’s theorem [48] the set \mathcal{Q} is nonempty if the marginals μ_1, \dots, μ_d are in convex order.

(c) Optimal transport with additional constraints:

$$\begin{aligned} \mathcal{Q} &= \{ \mu \in \Pi(\mu_1, \dots, \mu_d) : \int g_j d\mu = c_j \text{ for all } j = 1, \dots, N \} \\ \mathcal{H} &= \{ h \in C_b(\mathbb{R}^d) : h(x_1, \dots, x_d) = \sum_{i=1}^d h_i(x_i) + \sum_{j=1}^N \lambda_j (g_j(x_1, \dots, x_d) - c_j) \} \\ &\quad \text{for some } h_i \in C_b(\mathbb{R}), \lambda_j \in \mathbb{R} \} \end{aligned}$$

for some $g_1, \dots, g_N \in C_b(\mathbb{R}^d)$ and $c_1, \dots, c_N \in \mathbb{R}$. For related problems we refer to [6] and the references therein.

2.1 Regularization of the Superhedging Functional by Penalization

Our goal is to regularize the superhedging functional ϕ by considering the convolution

$$\begin{aligned} \phi_{\theta, \gamma}(f) &:= \inf_{h \in C_b(\mathcal{X})} \{ \phi(h) + \psi_{\theta, \gamma}(f - h) \} \\ &= \inf_{h \in \mathcal{H}} \left\{ \int h d\mu_0 + \int \beta_\gamma(f - h) d\theta \right\} \end{aligned} \tag{2.3}$$

where $\psi_{\theta, \gamma}(f) := \int \beta_\gamma(f) d\theta$ for a *sampling measure* $\theta \in \mathcal{P}(\mathcal{X})$, and $\beta_\gamma(x) := \frac{1}{\gamma} \beta(\gamma x)$ is a *penalty function* which is parametrized by $\gamma > 0$. We assume that $\beta : \mathbb{R} \rightarrow \mathbb{R}_+$ is a differentiable nondecreasing convex function such that $\lim_{x \rightarrow \infty} \beta(x)/x = \infty$. Its convex conjugate

$$\beta_\gamma^*(y) := \sup_{x \in \mathbb{R}} \{ xy - \beta_\gamma(x) \} \text{ for all } y \in \mathbb{R}_+,$$

satisfies $\beta_\gamma^*(y) = \beta^*(y)/\gamma$. Common examples are

- (a) the exponential penalty function $\beta(x) = \exp(x - 1)$ with conjugate $\beta^*(y) = y \log(y)$,
- (b) the L^p penalty function $\beta(x) = \frac{1}{p}(\max\{0, x\})^p$ with conjugate $\beta^*(y) = \frac{1}{q}y^q$ where $q = \frac{p}{p-1}$ for some $p > 1$.

In case that $\mathcal{H} = \mathbb{R}$ the functional (2.3) is a so-called *optimized certainty equivalent*, see Ben-Tal and Teboulle [8]. In the following result we show the dual representation of the regularized superhedging functional $\phi_{\theta, \gamma}$ and its convergence to ϕ .

Theorem 2.2 *Let $f \in C_b(\mathcal{X})$. Suppose there exists $\pi \in \mathcal{Q}$ such that $\pi \ll \theta$ and $\int \beta^*\left(\frac{d\pi}{d\theta}\right) d\theta < \infty$. Then*

$$\phi_{\theta, \gamma}(f) = \max_{\mu \in \mathcal{Q}} \left\{ \int f d\mu - \frac{1}{\gamma} \int \beta^*\left(\frac{d\mu}{d\theta}\right) d\theta \right\}. \tag{2.4}$$

Moreover,

$$\phi_{\theta, \gamma}(f) - \frac{\beta(0)}{\gamma} \leq \phi(f) \leq \phi_{\theta, \gamma}(f) + \frac{1}{\gamma} \int \beta^*\left(\frac{d\mu_\varepsilon}{d\theta}\right) d\theta + \varepsilon \tag{2.5}$$

whenever $\mu_\varepsilon \in \mathcal{Q}$ is an ε -optimizer of (2.2) such that $\mu_\varepsilon \ll \theta$ and $\int \beta^*\left(\frac{d\mu_\varepsilon}{d\theta}\right) d\theta < \infty$.

If $\hat{h} \in \mathcal{H}$ is a minimizer of (2.3) then $\hat{\mu} \in \mathcal{P}(\mathcal{X})$ defined by

$$\frac{d\hat{\mu}}{d\theta} := \beta'_\gamma(f - \hat{h}) \tag{2.6}$$

is a maximizer of (2.4).

2.2 Approximation of the Superhedging Functional

In this subsection we consider a sequence $\mathcal{H}^1 \subseteq \mathcal{H}^2 \subseteq \dots$ of subsets of \mathcal{H} , and set $\mathcal{H}^\infty := \bigcup_{m \in \mathbb{N}} \mathcal{H}^m$. For each $m \in \mathbb{N} \cup \{+\infty\}$, we define the *approximated superhedging functional* by

$$\phi^m(f) := \inf \left\{ \int h d\mu_0 : h \geq f \text{ for some } h \in \mathcal{H}^m \right\}. \tag{2.7}$$

For the approximation of $\phi(f)$ by $\phi^m(f)$, we need the following density condition on \mathcal{H}^∞ .

Condition (D) For every $\varepsilon > 0$ and $\mu \in \mathcal{P}(\mathcal{X})$ holds

- (a) for every $h \in \mathcal{H}$ there exists $h' \in \mathcal{H}^\infty$ such that $\int |h - h'| d\mu \leq \varepsilon$,
- (b) there exists $h'' \in \mathcal{H}^\infty$ such that $1_{K^c} \leq h''$ and $\int h'' d\mu \leq \varepsilon$ for some compact subset K of \mathcal{X} .

In Sect. 3 we will discuss Condition (D) in the context of multilayer feedforward networks. The condition allows for the following approximation result.

Proposition 2.3 *Assume that \mathcal{H}^∞ is a linear space which contains the constants. Under Condition (D) one has*

$$\lim_{m \rightarrow \infty} \phi^m(f) = \phi^\infty(f) = \phi(f)$$

for all $f \in C_b(\mathcal{X})$.

Given a sampling measure θ and a parametrized penalty function β_γ as in the previous subsection, we define the approximated version of the regularized superhedging functional by

$$\phi_{\theta,\gamma}^m(f) = \inf_{h \in \mathcal{H}^m} \left\{ \int h \, d\mu_0 + \int \beta_\gamma(f - h) \, d\theta \right\} \tag{2.8}$$

for all $f \in C_b(\mathcal{X})$. As a consequence of the two approximative steps $\phi_{\theta,\gamma}(f) \rightarrow \phi(f)$ for $\gamma \rightarrow \infty$ in Theorem 2.2 and $\phi^m(f) \rightarrow \phi(f)$ for $m \rightarrow \infty$ in Proposition 2.3 we get the following convergence result.

Proposition 2.4 *Suppose that \mathcal{H}^∞ satisfies Condition (D) and for every $\varepsilon > 0$ there exists an ε -optimizer μ_ε of (2.4) such that $\mu_\varepsilon \ll \theta$ and $\int \beta^*\left(\frac{d\mu_\varepsilon}{d\theta}\right) \, d\theta < +\infty$. Then, for every $f \in C_b(\mathcal{X})$ one has $\phi_{\theta,\gamma}^m(f) \rightarrow \phi(f)$ for $\min\{m, \gamma\} \rightarrow \infty$.*

The existence of such ε -optimizers as required in Theorem 2.2 and Proposition 2.4 is for example established in [12] in the context of multi-marginal optimal transport problems in \mathbb{R}^d with absolutely continuous marginals. In general, the existence of such ε -optimizers crucially depends on the choice of θ , see also Example 3.6 for a simple illustration.

3 Modelling Finite Dimensional Subspaces with Multilayer Feedforward Networks

This section explains the specific choice of approximative subspaces as built by neural networks. Generally, a feasible alternative to neural networks is to build these spaces via basis functions, like polynomials, which is for example pursued in [32] in the context of martingale optimal transport. In contrast to a basis approach, where functions are represented as a weighted sum over fixed basis functions, neural networks rely on the composition of layers of simple functions. This has shown to be an efficient way to approximate a large class of functions with relatively few parameters. Before going into the results, we give the required notation for neural networks.

3.1 Notation

The type of neural networks we consider are fully connected feed-forward neural networks. Those are mappings of the form

$$\mathbb{R}^d \ni x \mapsto A_l \circ \underbrace{\varphi \circ A_{l-1}}_{(l-1)\text{. layer}} \circ \cdots \circ \underbrace{\varphi \circ A_0}_{1\text{. layer}}(x)$$

where A_i are affine transformations and $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ is a nonlinear *activation* function that is applied elementwise, i.e. $\varphi((x_1, \dots, x_n)) = (\varphi(x_1), \dots, \varphi(x_n))$ for $(x_1, \dots, x_n) \in \mathbb{R}^n$.

Regarding dimensions, there is an input dimension $d \in \mathbb{N}$ and a hidden dimension $m \in \mathbb{N}$. This means A_0 maps from \mathbb{R}^d to \mathbb{R}^m , A_1, \dots, A_{l-1} map from \mathbb{R}^m to \mathbb{R}^m , and A_l maps from \mathbb{R}^m to \mathbb{R} . Each affine transformation A_j can trivially be represented as $A_j(x) = M_j x + b_j$ for a matrix M_j and a vector b_j . All these matrices and vectors together are the parameters of the network, which can be regarded as an element of \mathbb{R}^D for some $D \in \mathbb{N}$.

We will require the sets which contain all feed-forward neural networks with fixed structure (i.e. fixed number of layers and fixed dimensions) but unspecified parameter values. We denote by $\Xi \subset \mathbb{R}^D$ the sets of possible parameters for a fixed network structure (where formally, D depends on the structure of the network), and by $N_{l,d,m}(\xi) = A_l \circ \varphi \circ A_{l-1} \circ \dots \circ \varphi \circ A_0$ a particular neural network with l layers, input dimension d , hidden dimension m and parameters $\xi \in \Xi$. We denote the set of all such networks $N_{l,d,m}(\xi)$ for $\xi \in \Xi$ by $\mathfrak{N}_{l,d,m}(\Xi)$.

In the remainder of this section, we work with a fixed number of layers and input dimension, but allow for growing hidden dimension. For different hidden dimensions m , denote by Ξ_m the corresponding parameter sets. We define

$$\mathfrak{N}_{l,d} := \bigcup_{m \in \mathbb{N}} \mathfrak{N}_{l,d,m}(\Xi_m).$$

We want this definition to be independent of the precise choices of the parameter sets, which is why we make the standing assumption that the sets $\mathfrak{N}_{l,d,m}(\Xi_m)$ are growing in m . One way to make this explicit is:

Assumption 3.1 For any $l, d \in \mathbb{N}$ and a sequence of parameter sets Ξ_1, Ξ_2, \dots , where Ξ_m is regarded as a subset of \mathbb{R}^{D_m} for some $D_m \in \mathbb{N}$, we will always assume that $[-m, m]^{D_m} \subseteq \Xi_m$ and $\mathfrak{N}_{l,d,m}(\Xi_m) \subset \mathfrak{N}_{l,d,m+1}(\Xi_{m+1})$ for all $m \in \mathbb{N}$.

The only reason why we do not just set $\Xi_m \equiv \mathbb{R}^{D_m}$ is that in Proposition 3.7 we make the assumption of compact parameter sets. Further, we assume

Assumption 3.2 The activation function φ is continuous, nondecreasing and satisfies the limit properties $\lim_{x \rightarrow -\infty} \varphi(x) = 0$ and $\lim_{x \rightarrow +\infty} \varphi(x) = 1$.

3.2 Modelling \mathcal{H}^m Via Neural Networks

In the following we assume that \mathcal{H} is of the form

$$\mathcal{H} = \left\{ \sum_{j=1}^J e_j h_j \circ \pi_j + a : h_j \in C_b(\mathbb{R}^{d_j}), a \in \mathbb{R} \right\},$$

where $e_j \in C_b(\mathcal{X})$ and $\pi_j : \mathcal{X} \rightarrow \mathbb{R}^{d_j}$ are continuous functions for all $j = 1, \dots, J$. This form of \mathcal{H} includes many different problems, for instance the ones considered

in Example 2.1 (e.g. in (a) one has $\mathcal{H} = \{\sum_{j=1}^d h_j \circ \text{pr}_k : h_j \in C_b(\mathbb{R})\}$ where $\text{pr}_j(x) := x_j$ denotes the projection on the j -th marginal component).

We approximate \mathcal{H} by

$$\mathcal{H}^\infty = \left\{ \sum_{j=1}^J e_j h_j \circ \pi_j + a : h_j \in \mathfrak{N}_{l_j, d_j}, a \in \mathbb{R} \right\},$$

and its subspaces

$$\mathcal{H}^m = \left\{ \sum_{j=1}^J e_j h_j \circ \pi_j + a : h_j \in \mathfrak{N}_{l_j, d_j, m}(\Xi_{j, m}), a \in \mathbb{R} \right\}.$$

In this context the problems $\phi_{\theta, \gamma}^m(f)$ are given by

$$\begin{aligned} \phi_{\theta, \gamma}^m(f) &= \inf_{h \in \mathcal{H}^m} \left\{ \int h \, d\mu_0 + \int \beta_\gamma(f - h) \, d\theta \right\} \\ &= \inf_{a \in \mathbb{R}} \inf_{h_j \in \mathfrak{N}_{l_j, d_j, m}(\Xi_{j, m})} \left\{ \int \sum_{j=1}^J e_j h_j \circ \pi_j \, d\mu_0 + a \right. \\ &\quad \left. + \int \beta_\gamma \left(f - \sum_{j=1}^J e_j h_j \circ \pi_j - a \right) \, d\theta \right\} \\ &= \inf_{a \in \mathbb{R}} \inf_{\xi_j \in \Xi_{j, m}} \left\{ \int \sum_{j=1}^J e_j N_{l_j, d_j, m}(\xi_j) \circ \pi_j \, d\mu_0 + a \right. \\ &\quad \left. + \int \beta_\gamma \left(f - \sum_{j=1}^J e_j N_{l_j, d_j, m}(\xi_j) \circ \pi_j - a \right) \, d\theta \right\} \end{aligned}$$

for all $f \in C_b(\mathcal{X})$. The final formulation illustrates that the problem $\phi_{\theta, \gamma}^m(f)$ is now reduced to a finite dimensional problem of finding the optimal parameters in a neural network. Further, the overall objective depends smoothly on the parameters, and the parameters are unconstrained. In short, problem $\phi_{\theta, \gamma}^m(f)$ fits into the framework of machine learning problems that can be numerically solved by standard stochastic gradient descent based methods.

Under the standing Assumptions 3.1 and 3.2, the following lemma establishes situations when Condition (D), which is required for Proposition 2.3, is satisfied in the neural network setting.

Lemma 3.3 (a) \mathcal{H}^∞ satisfies the first part of Condition (D).

(b) If $\mathcal{X} = \mathbb{R}^d = \mathbb{R}^{d_1} \times \dots \times \mathbb{R}^{d_{J_0}}$ and $\pi_j = \text{pr}_j$, $e_j = 1$ for $j = 1, \dots, J_0 \leq J$, where pr_j is the projection from \mathbb{R}^d to j -th marginal component \mathbb{R}^{d_j} , then \mathcal{H}^∞

satisfies the second part of Condition (D). Further, the second part of Condition (D) is trivially satisfied whenever \mathcal{X} is compact.

Notably, part (b) can be seen as a large, but still exemplary case. Intuitively, the second part of Condition (D) is satisfied whenever the space \mathcal{H}^∞ is rich enough.

Remark 3.4 Later in the numerics we will usually work with a ReLU activation function, i.e. $\varphi(x) = \max\{0, x\}$. While this does not satisfy the latter limit property of Assumption 3.2, this is easily amendable: Basically, throughout the whole theory the assumptions will only be used to guarantee existence of neural networks with certain properties. Given Assumption 3.2, we will only require two layers ($l = 1$) to obtain the necessary results. In the numerics however, we use more layers. If more layers are given, one can also bundle several layers and regard them as one layer, with a different activation function. For example:

$$A_l \circ \varphi \circ \underbrace{A_{l-1} \circ \dots \circ A_1 \circ \varphi \circ A_0}_{\bar{\varphi}}$$

Whenever $\bar{\varphi}$ is a mapping of the form $(x_1, \dots, x_m) \mapsto (\bar{\varphi}(x_1), \dots, \bar{\varphi}(x_m))$, an $(l + 1)$ -layer network with activation function φ can represent any function that a two layer network with activation function $\bar{\varphi}$ can represent. For $\varphi(x) = \max\{0, x\}$ one can easily see that $\bar{\varphi}(x) = \min\{1, \max\{0, x\}\}$ is feasible, which satisfies Assumption 3.2.

3.3 Convergence

In this section we study in what sense $\phi_{\theta, \gamma}^m(f)$ converges to $\phi(f)$ for the approximation by neural networks.

First, we study the case of uniform convergence in m and γ , i.e. conditions for the convergence $\phi_{\theta, \gamma}^m(f) \rightarrow \phi(f)$ for $\min\{m, \gamma\} \rightarrow \infty$. This is subject of Remark 3.5 below, which is a summary of results established in Sects. 2 and 3.2. The two approximative steps leading to uniform convergence are $\phi_{\theta, \gamma}(f) \rightarrow \phi(f)$ for $\gamma \rightarrow \infty$ and $\phi^m(f) \rightarrow \phi(f)$ for $m \rightarrow \infty$.

On the other hand, sometimes the convergence $\phi_{\theta, \gamma}(f) \rightarrow \phi(f)$ for $\gamma \rightarrow \infty$ is not satisfied even though practically one obtains a good approximation. One such case is given in Example 3.6. Even if uniform convergence does not hold, one can still often connect problems $\phi_{\theta, \gamma}^m(f)$ and $\phi(f)$. This is done by the approximative steps $\phi^m(f) \rightarrow \phi(f)$ for $m \rightarrow \infty$ and $\phi_{\theta, \gamma}^m(f) \rightarrow \phi^m(f)$ for $\gamma \rightarrow \infty$, where the latter is subject of Proposition 3.7. Here, instead of the strong assumption required for $\phi_{\theta, \gamma}(f) \rightarrow \phi(f)$, the convergence $\phi_{\theta, \gamma}^m(f) \rightarrow \phi^m(f)$ can be shown by assuming that all parameter sets of the neural networks are compact.

Remark 3.5 Under the assumptions of Lemma 3.3, Proposition 2.3 implies $\phi^m(f) \rightarrow \phi(f)$ for $m \rightarrow \infty$. Further, given the existence of ε -optimizers for every $\varepsilon > 0$ as required in Theorem 2.2, convergence $\phi_{\theta, \gamma}(f) \rightarrow \phi(f)$ for $\gamma \rightarrow \infty$ holds. Given both assumptions, Proposition 2.4 yields $\phi_{\theta, \gamma}^m(f) \rightarrow \phi(f)$ for $\min\{m, \gamma\} \rightarrow \infty$. The convergence $\phi_{\theta, \gamma}^m(f) \rightarrow \phi_{\theta, \gamma}(f)$ for $m \rightarrow \infty$ is a trivial consequence.

Example 3.6 Let $\mathcal{X} = [0, 1]^2$, $\mu_1 = \mu_2 = \delta_0$ and $f(x_1, x_2) = -|x_1 - x_2|$. Let $\mathcal{Q} = \Pi(\mu_1, \mu_2)$ be the set of all measures in \mathcal{X} with first marginal μ_1 and second marginal μ_2 , so that

$$\phi(f) = \sup_{\mu \in \Pi(\mu_1, \mu_2)} \int f d\mu$$

Obviously, $\phi(f) = f(0, 0) = 0$. Note that $\mathcal{Q} = \{\mu_1 \otimes \mu_2\}$ so that $\mu_0 = \mu_1 \otimes \mu_2 = \delta_{(0,0)}$.

Consider two possible reference measures, $\theta^{(1)} = \mathcal{U}([0, 1]^2)$ being the uniform distribution on $[0, 1]^2$, and $\theta^{(2)} = \mu_1 \otimes \mu_2 = \delta_{(0,0)}$. For $\theta^{(2)}$ it is obvious that the existence of ε -optimizers as required in Theorem 2.2 is given, since $\theta^{(2)}$ itself is the optimizer of $\phi(f)$. Hence $\phi_{\theta^{(2)}, \gamma}(f) \rightarrow \phi(f)$ for $\gamma \rightarrow \infty$ holds.

On the other hand, there does not exist $\nu \in \Pi(\mu_1, \mu_2)$ with $\nu \ll \theta^{(1)}$, and hence $\phi_{\theta^{(1)}, \gamma}(f) = -\infty$. However, by first approximating $\phi(f)$ by $\phi^m(f)$, the functional becomes smoother: Roughly speaking, the marginal constraints are slightly relaxed. This becomes obvious when studying the dual formulations

$$\begin{aligned} \phi_{\theta^{(1)}, \gamma}(f) &= \inf_{h_1, h_2 \in C_b([0, 1])} \left\{ h_1(0) + h_2(0) + \int \beta_\gamma(f - h_1 - h_2) d\theta^{(1)} \right\} \\ \phi_{\theta^{(1)}, \gamma}^m(f) &= \inf_{h_1, h_2 \in \mathfrak{N}_{l, 1, m}(\Xi_m)} \left\{ h_1(0) + h_2(0) + \int \beta_\gamma(f - h_1 - h_2) d\theta^{(1)} \right\} \end{aligned}$$

While one easily finds sequences of functions in $C_b([0, 1])$ so that the values at 0 go to minus infinity but the penalty term stays bounded, this is impossible with functions in $\mathfrak{N}_{l, 1, m}(\Xi_m)$ given that the activation function is continuous and the parameter sets are compact. So there is hope to establish the convergence $\phi_{\theta^{(1)}, \gamma}^m(f) \rightarrow \phi^m(f)$ for $\gamma \rightarrow \infty$, which will indeed be a consequence of the following result.

Proposition 3.7 Fix $m \in \mathbb{N}$. Given that all parameter sets $\Xi_{j, m}$ for $j = 1, \dots, J$ of the neural networks occurring in \mathcal{H}^m are compact and θ is strictly positive (i.e. θ gives positive mass to every non-empty open set), it holds $\phi_{\theta, \gamma}^m(f) \rightarrow \phi^m(f)$ for $\gamma \rightarrow \infty$.

4 Numerical Examples

This section aims at showcasing how various frequently studied problems that fall into the theoretical framework of the previous sections can be implemented simply and effectively with neural networks. The examples focus on toy problems that allow an objective evaluation of the numerical results and give the reader an idea about the strengths and weaknesses of the presented approach. We chose a very basic implementation using Tensorflow and the Adam optimizer.² As for the network architecture: In all the examples, \mathcal{H} is as described in Sect. 3, and $\mathfrak{N}_{l, k, m}$ always approximates $C_b(\mathbb{R}^d)$. To approximate $C_b(\mathbb{R}^d)$ we use a five layer ($l = 4$ in the previous chapter)

² All used code is available on <https://github.com/stephaneckstein/transport-and-related>.

ReLU-network with hidden dimension $64 \cdot d$. We did not perform a hyper parameter search to obtain this architecture, but rather oriented ourselves at papers with comparable settings (e.g. at [29,45,51]). Notably, increasing the complexity (number of layers or hidden dimension) further did not change the numerical results significantly in the cases tested, so we believe the structure chosen to be adequate for the problems considered.

Simply put, the implementation works as follows: We perform a normal stochastic gradient type optimization (outsourced to the Adam optimizer) for a certain number of iterations to find near optimal parameters of the network. At each iteration during this process, the expectations in the objective function are replaced by averages over a fixed number (called *batch size*) of random points from the respective distributions. To obtain the numerical approximation $\hat{\phi}_{\theta,\gamma}^m(f)$ of $\phi_{\theta,\gamma}^m(f)$, we finally average the sample objective values over the last roughly 5% of iterations. This is referred to as the dual value. Alternatively, one can use formula (2.6) to obtain sample points from an approximate optimizer ν^* of the primal problem and numerically evaluate $\int f d\nu^*$, which is referred to as the primal value (more details on how to work with such an approximative optimizer ν^* is given in Sect. 4.5). If not stated otherwise, all reported values are dual values.

The numerical procedure we use can likely be improved by fine-tuning parameters or by using more complex network architectures. For example batch normalization is applied in a related setting in [13] which appears to significantly speed up the optimization.

4.1 Optimal Transport and Fréchet–Hoeffding Bounds

With this first problem, we study the effects of different penalty functions, penalty factor, batch size and number of iterations of Let $\mathcal{X} = [0, 1]^d$, $\theta = \mathcal{U}([0, 1]^d)$ (where $\mathcal{U}(\cdot)$ denotes the uniform distribution) and $\mathcal{Q} = \{\nu \in \mathcal{P}(\mathcal{X}) : \nu_i = \mathcal{U}([0, 1])\}$, where ν_i is the i -th marginal of ν . For some fixed $z \in [0, 1]^d$, define the function $f : [0, 1]^d \rightarrow \mathbb{R}_+$ by³

$$f(x) = \begin{cases} 1, & \text{if } x_i \leq z_i \text{ for all } i \in \{1, 2, \dots, d\}, \\ 0, & \text{else.} \end{cases}$$

The value $\phi(f) = \sup_{\nu \in \mathcal{Q}} \int f d\nu$ corresponds to the maximum value of a d -dimensional copula at point z . By the Fréchet-Hoeffding bounds we have an analytical solution to this problem, which is

$$\phi(f) = \min_{i \in \{1, \dots, d\}} z_i.$$

In Fig. 2 we observe how $\hat{\phi}_{\theta,\gamma}^m(f)$ depends on the number of iterations of the Adam optimizer and the batch size. We observe that while higher batch sizes lead

³ The function f here is not continuous. Since the optimal transport problem is continuous from below (see e.g. [36]), the representation (2.2) nevertheless holds for all bounded measurable functions f .




Fig. 2 Fréchet-Hoeffding bounds: $d = 2, z_1 = 0.5, z_2 = 0.75$. Comparison of L^2 penalty function $\beta_\gamma(x) = \gamma \max\{0, x\}^2$ and exponential penalty function $\beta_\gamma(x) = \frac{\exp(\gamma x - 1)}{\gamma}$. The values plotted are running averages over the last 1000 iterations. The dotted red line is the true value $\phi(f)$. The dotted blue lines are bounds from below for $\phi_{\theta, \gamma}(f)$ obtained by Eq. (2.5) in Theorem 2.2 for the respective choices of γ (Color figure online)

to more stable convergence, the speed of convergence appears not strongly related to batch size. This suggests that increasing batch sizes might lead to both quick and finally stable performance.⁴ Since L^2 penalization appears more stable, we will mostly use this penalization for the rest of the applications. Further, the figure illustrates that the numerical solutions appear to approximately obtain the lower bounds for $\phi_{\theta, \gamma}(f)$ as given by Eq. (2.5) in Theorem 2.2. I.e. one approximately has $\phi(f) \approx \phi_{\theta, \gamma}^m(f) + \frac{1}{\gamma} \int \beta^*(\frac{d\hat{\mu}}{d\theta})d\theta$ where $\hat{\mu}$ is an optimizer of $\phi(f)$.⁵

4.2 Multi-marginal Optimal Transport

The aim of this example is to compare the approach of this paper with existing methods for a numerically challenging problem. Let $\mathcal{X} = (\mathbb{R}^D)^M$, where M denotes the number of marginals and D denotes the dimension of each marginal. Let μ_i for $i = 1, \dots, M$ be K -mixtures of normal distributions with randomly chosen parameters, and define $\mathcal{Q} = \Pi(\mu_1, \dots, \mu_M)$. For $p, q \geq 1$ let

$$f(x) := -\left(\sum_{j=1}^D \left|\sum_{i=1}^M (-1)^i x_{i,j}\right|^q\right)^{p/q},$$

where we write $x = (x_{i,j}) \in \mathcal{X}$ with $i = 1, \dots, M, j = 1, \dots, D$. Note that for two marginals, one has $-\phi(f) = W_{p,q}^p(\mu_1, \mu_2)$, where $W_{p,q}$ is the Wasserstein- p distance with L^q norm on \mathbb{R}^d .

In Table 1 we compare optimal values to this problem arising from different algorithmic approaches. We compare linear programming methods based on discretization

⁴ See also [46] and references therein for related concepts on how to optimally tune the optimization procedure. In this paper however, we decided to stick with standard parameters of the Adam optimizer and fixed batch size. This is done to avoid another layer of complexity when evaluating the numerical results.

⁵ Here, $\hat{\mu}$ is chosen as the optimizer which is uniform on the cubes $[0, z]$ and $[z, 1]$.

Table 1 Multi-marginal optimal transport: Numerical values for $-\phi(f)$ arising from different numerical schemes

(M, D, K)	LP		NN		RKHS	Ref
	MC	Quantization	Dual	Primal	Laplace	com.
p = q = 2						
(2, 1, 1)	0.403 (0.084)	0.408 (0.026)	0.413	0.401	0.364 (0.006)	0.405
(2, 1, 6)	3.337 (0.320)	3.263 (0.115)	3.279	3.258	2.444 (0.018)	3.269
(5, 2, 6)	8.978 (8.233)	3.073 (0.231)	3.123	3.041	DNC	–
p = 1, q = 2						
(2, 1, 6)	1.536 (0.071)	1.537 (0.025)	1.537	1.531	1.471 (0.009)	1.533
(5, 2, 6)	2.845 (1.314)	1.741 (0.064)	1.753	1.740	DNC	–
(10, 3, 6)	10.235 (3.576)	6.744 (0.074)	6.759	6.743	DNC	–
$\tilde{f}(x) = f(x) \cdot \sin(\sum_{i=1}^M x_i, 1)$						
(5, 2, 6)	16.814 (0.893)	17.380 (0.043)	18.001	17.539	DNC	–
(10, 3, 6)	24.618 (2.332)	23.615 (0.107)	34.235	32.521	DNC	–

The numbers in brackets are empirical standard deviations over 100 runs. LP denotes linear programming, based on either sampling the marginals randomly (MC) or using the quantization approach from [41, Algorithm 4.5] to approximate marginals (quantization). The neural network (NN) implementation is based on L^2 penalization with $\gamma = M \cdot D \cdot 500$. For the reproducing kernel Hilbert space solution (RKHS) as described in [27, Algorithm 3] we use a Laplace kernel and the same penalization as for the NN-method. For the final two rows, we report numerical values for $-\phi(\tilde{f})$. DNC entries did not converge. The final column are analytical reference values given by the comonotone coupling for two marginals

of the marginals, the neural network method presented in this paper, and a reproducing kernel Hilbert space (RKHS) approach as described in [27, Algorithm 3]. For the linear programming methods, we use a maximal number of variables of 10^6 , which was around the boundary so that Gurobi [31] was still able to solve the resulting linear program on our computer. Regarding the RKHS algorithm we have to mention that it is the only method that is not building on an established package like Tensorflow or Gurobi. Hence efficiency with respect to running time and tuning of hyperparameters might be far from optimal for this approach. Notably, switching from exponential penalization as used in [27] to L^2 penalization was already a slight improvement. For the precise specifications of each algorithm and the problem setting, we refer to the code on https://github.com/stephaneckstein/OT_Comparison.

Evaluating the results, we find that the neural network based method and the linear programming method with quantization appear to work best. Surprisingly, even for the

case with 10 marginals (where the linear program can only use 4 points to approximate each marginal!), the quantization method achieved a similar value as the neural network method for $-\phi(f)$. We believe the reason is that the function f is very smooth which a quantization approach can exploit. Hence we slightly changed f to \tilde{f} in the final two test cases, which makes the function less regular. These are the only cases where the neural network solution and the quantization method strongly differ. In the final case, the quantization approach still has to approximate each marginal distribution with just 4 points, while the neural network method can use millions of points. From this standpoint, one can place higher trust in the neural network solution, even though we have no analytical reference value to compare it against.

Initially, we included a fourth method based on the approach in this paper, but with a polynomial basis instead of neural networks. This performed very badly however (at least when using a standard monomial basis), and hence we omitted the results in this table.

4.3 Martingale Optimal Transport

In martingale optimal transport, the optimal transport problem is extended by imposing a martingale constraint on top of marginal constraints. Dimensions are regarded as discrete time-steps and the measures in \mathcal{Q} are distributions of discrete stochastic processes $(X_t)_{t=1,\dots,d}$ with fixed marginal distributions as well as the condition that the process is a martingale.

Here, we consider a simple example with $d = 2$, where an analytical solution is known. This example is taken from [2]. Let $\mathcal{X} := [-1, 1] \times [-2, 2]$, $\theta := \mathcal{U}(\mathcal{X})$ and set

$$\mathcal{Q} := \left\{ \nu = \nu_1 \otimes K : \nu_1 = \mathcal{U}([-1, 1]), \nu_2 = \mathcal{U}([-2, 2]), \right. \\ \left. x = \int_{-2}^2 y K(x, dy) \text{ holds } \nu_1\text{-a.s.} \right\}.$$

For $f = -|x - y|^\rho$ one gets $\phi(f) = -1$ for all $\rho > 2$. We implement this problem with $\rho = 2.3$, where we use the L^2 penalty function for different values of γ . The results are shown in Fig. 3. One can see that while for values of γ up to around 1280, the behavior of the optimal value is approximately as predicted by Eq. (2.5) in Theorem 2.2, in that the error decreases by roughly a factor of two if γ is increased by a factor of two. For larger values of γ however, numerical instabilities occur and the optimizer cannot find the true optimum. This is indicated by the fact that the value $\hat{\phi}_{\theta,\gamma}^m(f)$ is above $\phi(f) = -1$.

4.4 Portfolio Optimization

Consider a market with two assets, where the distribution of returns for each individual asset is given, but not the joint distribution. An investor wants to maximize his or her worst-case utility from investing into the two assets. Here, the utility of the investor




Fig. 3 Martingale optimal transport: Mean numerical optimal values and 95% confidence bounds over 100 independent runs for different values of γ (L^2 penalization). The network is trained for 20,000 iterations with batch size 1024. The true optimal value of the unpenalized problem is -1

is characterized by a mean-variance objective. While the mean is fully characterized by the marginal distributions, the worst case considers all possible variances of the portfolio, which depend on the joint distribution of the assets.

The following example is taken from [42]: Let $\mathcal{X} = [0, 1] \times [0, 2]$. Let $\theta_1 = \mathcal{U}([0, 1])$ and $\theta_2 = \mathcal{U}([0, 1]) \circ \varphi^{-1}$, where $\varphi(x) = 2x^2$. Let $\mathcal{Q} = \{v \in \mathcal{P}(\mathcal{X}) : v_1 = \theta_1, v_2 = \theta_2\}$. We will solve the following robust mean-variance portfolio optimization problem

$$\begin{aligned} & \sup_{x \in [0,1]} -\phi(-f_x) \\ & := \sup_{x \in [0,1]} \inf_{v \in \mathcal{Q}} \int (1-x)\xi_1 + x\xi_2 \\ & \quad - \lambda \left((1-x)\xi_1 + x\xi_2 - (1-x) \int_0^1 \zeta_1 \theta_1(d\zeta_1) - x \int_0^2 \zeta_2 \theta_2(d\zeta_2) \right)^2 v(d\xi). \end{aligned}$$

where $\lambda \geq 0$ is the risk aversion. The integral over the term inside the large brackets is the variance of the portfolio. For the analytical solution, see Example 1 of [42]. We implemented the above problem in two ways. For the first, we choose the reference measure $\theta^{(1)} = \theta_1 \otimes \theta_2$. For the second, we use the reference measure $\theta^{(2)} = 0.5\theta^{(1)} + 0.5(\mathcal{U}([0, 1]) \circ (\text{Id}, \varphi)^{-1})$,⁶ i.e. half the product measure, half the perfectly correlated measure. The second version may correspond to our intuition that the optimal coupling should include positive correlation. More precisely: The choice of reference measure θ always has the implicit objective to lead to narrow bounds in Eq. (2.5) in Theorem

⁶ Id denotes the identity mapping.




Fig. 4 Portfolio optimization under dependence uncertainty: As reference measure we take either the product measure or a positively correlated measure. We use L^2 penalization with $\gamma = 160$. The network is trained with batch size 2^{13} for 40,000 iterations

2.2. In this example, if one presumes that an optimal measure $\nu^* \in \mathcal{Q}$ has mass near the perfectly correlated diagonal, it makes sense to choose a reference measure which puts mass in this region, as does $\theta^{(2)}$. The results are reported in Fig. 4. As expected, the second version yields results closer to the analytical solution.

4.5 Bounds on the Distribution of a Sum of Dependent Random Variables

In this section, the objective is to find bounds for the probability $\mathbb{P}(X_1 + X_2 + \dots + X_d \geq s)$ for some $s \in \mathbb{R}$, where the individual distributions of X_i are known, but not their joint distribution. This problem is in strong relation to calculating worst- and best-case value at risks under dependence uncertainty, see also [19,21,43].

Let $\mathcal{X} = \mathbb{R}^d$. For given marginals $\mu_1, \dots, \mu_d \in \mathcal{P}(\mathbb{R})$ and $\mathcal{Q} = \Pi(\mu_1, \dots, \mu_d)$, the problem statement is

$$\phi(f) := \sup_{\nu \in \mathcal{Q}} \int 1_{\{x_1 + \dots + x_d \geq s\}} \nu(dx_1, \dots, dx_d).$$

For simplicity, we consider the case $d = 2$, and $\mu_i = \mathcal{U}([0, 1])$. Let $s = 1.9$. Every optimal measure $\nu \in \mathcal{Q}$ gives mass $1/10$ uniformly to the line section $\{(x, 1.9 - x) : 0.9 \leq x \leq 1\}$, while the rest of the mass is irrelevant as long as the marginal condition is satisfied. This leads to an optimal value $\phi(f) = 0.1$. For the natural choice of reference measure $\theta = \mathcal{U}([0, 1]^2)$, every optimal measure is singular with respect to θ , and thus one can expect high errors by penalization. An implementation with this reference measure and L^2 penalization with $\gamma = 320$ leads to $\hat{\phi}_{\theta, \gamma}^m(f) \approx 0.0881$.⁷

Updating the reference measure: To obtain a more accurate value, we make use of Eq. (2.6) in Theorem 2.2. Recall that an optimizer $\hat{\nu} \in \mathcal{Q}$ of $\phi_{\theta, \gamma}(f)$ is given by $\frac{d\hat{\nu}}{d\theta} = \beta'(f - \hat{h})$, where \hat{h} is the dual optimizer of $\phi_{\theta, \gamma}(f)$. Taking $\hat{\nu}$ as a reference measure

⁷ The network was trained for 20000 iterations with batch size 1024.




Fig. 5 Bounds on the distribution of the sum of dependent variables: Sampled points from the numerically optimal measure ν^* and the corresponding empirical marginal distributions

instead of θ can only reduce the error by penalization, since $\phi_{\hat{\nu},\gamma}(f) \geq \phi_{\theta,\gamma}(f)$ holds by convexity of β_γ^* . Implementing the problem with $\hat{\nu}$ as a reference measure has to be done approximately, since the true optimizer \hat{h} is unknown and replaced by the numerical optimal solution. We denote the numerically obtained optimal measure by ν^* .

To implement $\phi_{\nu^*,\gamma}^m(f)$ requires sampling points from ν^* . This is non-trivial since ν^* is only given by $\frac{d\nu^*}{d\theta}$. We implemented this by an acceptance-rejection method as described in [23]. This is very slow, as the number of rejections increases with the maximum of the Radon-Nikodym derivative. Sampling efficiently in such a situation is difficult, see e.g. [33] for an overview of existing methods and a proposed new one.

Figure 5 illustrates the optimal measure ν^* . The measure ν^* looks comparable to an optimal solution of $\phi(f)$, while simultaneously being driven towards the reference measure θ . One obtains $\phi_{\nu^*,\gamma}^m(f) \approx 0.0982$, which is close to the true optimal value 0.1.

In the following, we briefly discuss the rearrangement algorithm [21,43] which is tailored to this type of problem. In contrast to the presented approach, which relies on sampling from the involved marginal distributions, the rearrangement algorithm mainly relies on the (inverse of) the cumulative distribution function of the marginals. The rearrangement algorithm achieves similar or even better accuracy in higher dimensional settings and with different marginals (e.g. Pareto marginals).

The case of higher dimensions scales well with the approach taken here. However, the base time in low dimensions is higher than that of the rearrangement algorithm, and further heavy tailed marginals like the Pareto distribution can lead to less accuracy.⁸

4.6 Generative Adversarial Networks (GANs)

The objective in GANs is to create new sample points from a measure μ , of which only an empirical distribution $\tilde{\mu}$ is known (see e.g. [3,28]). Usually, the measure μ might refer to the uniform distribution over some very large set of images. The set $\tilde{\mu}$ is then just the uniform distribution over a small subset of these images. The goal is to sample new images that are not already present in the given subset, but that might plausibly have been samples from the measure μ .

To proceed, we first take some latent probability measure τ . The goal is to obtain a function G such that μ and the push-forward measure $\tau \circ G^{-1}$ are close (in a sense to be specified), and thus the pseudo samples for μ can be obtained by sampling from τ and applying G . To find such a function G out of a class of functions \mathcal{G} , one can only use $\tilde{\mu}$ instead of μ (thus μ does not enter the formal problem statement). The closeness of $\tilde{\mu}$ and $\tau \circ G^{-1}$ is measured by different distances in GANs. In the Wasserstein GAN, the first Wasserstein distance $W_1(\cdot, \cdot)$ (see e.g. [50]) is used, and the objective is

$$\arg \min_{G \in \mathcal{G}} W_1(\tau \circ G^{-1}, \tilde{\mu}).$$

The above can be generalized to arbitrary transport distances instead of the first Wasserstein distance. To put this into our setting, let \mathcal{G} a set of functions that map into \mathcal{X}_1 . Let $\mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_1$ and for $G \in \mathcal{G}$ define $\mathcal{Q}_G := \{v \in \mathcal{P}(\mathcal{X}) : v_1 = \tau \circ G^{-1}, v_2 = \tilde{\mu}\}$. For a cost function c , arbitrary transport type GANs can be expressed via

$$\arg \min_{G \in \mathcal{G}} -\phi_G(-c) = \arg \min_{G \in \mathcal{G}} \inf_{v \in \mathcal{Q}_G} \int c \, dv$$

If c is a metric, the above corresponds to the Wasserstein GAN.

Since it is difficult to objectively evaluate GAN setups, we omit numerical results in this section. The interested reader can see the code on GitHub, where the toy problems appearing in [29] are implemented for different functions c .⁹

⁸ We refer to <https://github.com/stephaneckstein/transport-and-related>, where different specifications (including Pareto marginals) of the problem in this section are implemented.

⁹ For the implementation, we merely adjusted code from https://github.com/igul222/Improved_wgan_training to our method. See again <https://github.com/stephaneckstein/transport-and-related>.

5 Proofs

5.1 Proof of Theorem 2.2

- (1) We first show (2.4) by verifying that $\phi_{\theta,\gamma}$ is real-valued and continuous from above on $C_b(\mathcal{X})$. To that end, as shown in Appendix A one has $\phi^*(\mu) = \sup_{h \in \mathcal{H}} (\int h d\mu - \int h d\mu_0)$ for all $\mu \in \mathcal{P}(\mathcal{X})$, so that

$$\mu \in \mathcal{Q} \text{ if and only if } \int h d\mu = \int h d\mu_0 \text{ for all } h \in \mathcal{H}. \tag{5.1}$$

Since $\beta_\gamma(x) \geq xy - \frac{1}{\gamma}\beta^*(y)$ for all $x \in \mathbb{R}$ and $y \in \mathbb{R}_+$, it follows that

$$\int \beta_\gamma(f - h) d\theta \geq \int f - h d\pi - \frac{1}{\gamma} \int \beta^*\left(\frac{d\pi}{d\theta}\right) d\theta$$

so that

$$\begin{aligned} \phi_{\theta,\gamma}(f) &= \inf_{h \in \mathcal{H}} \left\{ \int h d\pi + \int \beta_\gamma(f - h) d\theta \right\} \\ &\geq \int f d\pi - \frac{1}{\gamma} \int \beta^*\left(\frac{d\pi}{d\theta}\right) d\theta > -\infty \end{aligned}$$

for all $f \in C_b(\mathcal{X})$. This shows that $\phi_{\theta,\gamma}$ is real-valued on $C_b(\mathcal{X})$. Further, let (f^k) be a sequence in $C_b(\mathcal{X})$ such that $f^k \downarrow f$. For every $h \in \mathcal{H}$ the monotone convergence theorem implies $\int \beta_\gamma(f^k - h) d\theta \rightarrow \int \beta_\gamma(f - h) d\theta$, so that $\phi_{\theta,\gamma}(f^k) \downarrow \phi_{\theta,\gamma}(f)$. Hence, it follows from the nonlinear Daniell-Stone theorem (see Proposition A.1) that

$$\phi_{\theta,\gamma}(f) = \max_{\mu \in \mathcal{P}(\mathcal{X})} \left\{ \int f d\mu - \phi_{\theta,\gamma}^*(\mu) \right\} \text{ for all } f \in C_b(\mathcal{X}),$$

where the convex conjugate is given by

$$\phi_{\theta,\gamma}^*(\mu) = \phi^*(\mu) + \psi_{\theta,\gamma}^*(\mu) = \begin{cases} \frac{1}{\gamma} \int \beta^*\left(\frac{d\mu}{d\theta}\right) d\theta & \text{if } \mu \in \mathcal{Q} \text{ and } \mu \ll \theta \\ \infty & \text{else.} \end{cases}$$

Indeed, the convex conjugate of the convolution $\inf_{f \in C_b(\mathcal{X})} \{\phi(f) + \psi_{\theta,\gamma}(\cdot - f)\}$ is given as the sum of the convex conjugates ϕ^* and $\psi_{\theta,\gamma}^*$. By (5.1) one has $\phi^*(\mu) = 0$ if $\mu \in \mathcal{Q}$ and $\phi^*(\mu) = +\infty$ otherwise. Moreover,

$$\begin{aligned} \psi_{\theta,\gamma}^*(\mu) &= \sup_{f \in C_b(\mathcal{X})} \left\{ \int f d\mu - \int \beta_\gamma(f) d\theta \right\} \\ &= \sup_{f \in C_b(\mathcal{X})} \left\{ \int f \frac{d\mu}{d\theta} - \beta_\gamma(f) d\theta \right\} = \int \beta_\gamma^*\left(\frac{d\mu}{d\theta}\right) d\theta \end{aligned}$$

if $\mu \ll \theta$ and $\psi_{\theta,\gamma}^*(\mu) = +\infty$ otherwise.

(2) We next show (2.5). On the one hand, one has

$$\begin{aligned} \phi_{\theta,\gamma}(f) &= \inf_{h \in \mathcal{H}} \left\{ \int h \, d\mu_0 + \int \beta_\gamma(f - h) \, d\theta \right\} \\ &\leq \inf_{\substack{h \in \mathcal{H}: \\ h \geq f}} \int h \, d\mu_0 + \beta_\gamma(0) = \phi(f) + \frac{\beta(0)}{\gamma}. \end{aligned}$$

On the other hand, for every ε -optimizer $\mu_\varepsilon \in \mathcal{Q}$ of (2.2) such that $\mu_\varepsilon \ll \theta$ one has

$$\begin{aligned} \phi(f) &\leq \int f \, d\mu_\varepsilon + \varepsilon \leq \int f \, d\mu_\varepsilon - \phi_{\theta,\gamma}^*(\mu_\varepsilon) + \phi_{\theta,\gamma}^*(\mu_\varepsilon) \\ &+ \varepsilon \leq \phi_{\theta,\gamma}(f) + \frac{1}{\gamma} \int \beta^*\left(\frac{d\mu_\varepsilon}{d\theta}\right) \, d\theta + \varepsilon \end{aligned}$$

with the convention $-\infty + \infty = +\infty$.

(3) Let $\hat{h} \in \mathcal{H}$ be a minimizer of (2.3), i.e. $\phi_{\mu,\gamma}(f) = \int \hat{h} \, d\mu_0 + \int \beta_\gamma(f - \hat{h}) \, d\theta$. Defining $h_\lambda := \hat{h} + \lambda h$ for an arbitrary $h \in \mathcal{H}$, the first order condition

$$\left. \frac{d}{d\lambda} \right|_{\lambda=0} \left(\int h_\lambda \, d\mu_0 + \int \beta_\gamma(f - h_\lambda) \, d\theta \right) = 0$$

implies

$$\int h \, d\mu_0 - \int \beta'_\gamma(f - \hat{h})h \, d\theta = 0.$$

This shows that the probability measure $\hat{\mu}$ with Radon-Nikodým derivative $\frac{d\hat{\mu}}{d\theta} := \beta'_\gamma(f - \hat{h})$ satisfies $\int h \, d\mu_0 = \int h \, d\hat{\mu}$ for all $h \in \mathcal{H}$, which in view of (5.1) satisfies $\hat{\mu} \in \mathcal{Q}$. Integrating the identity $\beta_\gamma(x) = x\beta'_\gamma(x) - \beta_\gamma^*(\beta'_\gamma(x))$ with $x = f - \hat{h}$ w.r.t. θ , one obtains

$$\int \beta_\gamma(f - \hat{h}) \, d\theta = \int f - \hat{h} \, d\hat{\mu} - \int \beta_\gamma^*\left(\frac{d\hat{\mu}}{d\theta}\right) \, d\theta$$

which shows that

$$\phi_{\theta,\gamma}(f) = \int \hat{h} \, d\mu_0 + \int \beta_\gamma(f - \hat{h}) \, d\theta = \int f \, d\hat{\mu} - \int \beta_\gamma^*\left(\frac{d\hat{\mu}}{d\theta}\right) \, d\theta.$$

As a consequence, $\hat{\mu} \in \mathcal{Q}$ is a maximizer of (2.4). □

5.2 Proof of Proposition 2.3

Fix $f \in C_b(\mathcal{X})$. That $\lim_{m \rightarrow \infty} \phi^m(f) = \phi^\infty(f) \geq \phi(f)$ follows from the definition of \mathcal{H}^∞ . Moreover, for every $\varepsilon > 0$ the Condition (D) guarantees $h \in \mathcal{H}^\infty$ and $K \subseteq \mathcal{X}$

such that $1_{K^c} \leq h$ and $\int h d\mu_0 \leq \varepsilon$. Hence, $\phi^\infty(1_{K^c}) \leq \int h d\mu_0 \leq \varepsilon$ and Dini’s lemma implies that ϕ^∞ is continuous from above on $C_b(\mathcal{X})$. By Proposition A.1 it follows that

$$\phi^\infty(f) = \max_{\mu \in \mathcal{P}(\mathcal{X})} \left\{ \int f d\mu - \phi^{\infty*}(\mu) \right\}.$$

Similar to (A.2) its convex conjugate is given by

$$\phi^{\infty*}(\mu) = \sup_{h \in \mathcal{H}^\infty} \left(\int h d\mu - \int h d\mu_0 \right) \leq \sup_{h \in \mathcal{H}} \left(\int h d\mu - \int h d\mu_0 \right) = \phi^*(\mu).$$

It remains to show that for $h \in \mathcal{H}$ and $\mu \in \mathcal{P}(\mathcal{X})$ with $\int h d\mu - \int h d\mu_0 > 0$ there exists $h' \in \mathcal{H}^\infty$ such that $\int h' d\mu - \int h' d\mu_0 > 0$. But this follows directly from the first part of Condition (D) for the probability measure $\frac{1}{2}\mu + \frac{1}{2}\mu_0$. Indeed, there exists a sequence (h^n) in \mathcal{H}^∞ such that $h^n \rightarrow h$ in $L^1(\mu)$ and in $L^1(\mu_0)$, which shows that $\int h^n d\mu - \int h^n d\mu_0 > 0$ for n large enough. \square

5.3 Proof of Proposition 2.4

Observe that

$$\phi^m(f) + \beta_\gamma(0) \geq \inf_{\substack{h \in \mathcal{H}^m \\ h \geq f}} \left\{ \int h d\mu_0 + \int \beta_\gamma(f - h) d\theta \right\} \geq \phi_{\theta, \gamma}^m(f) \geq \phi_{\theta, \gamma}(f)$$

where the first inequality uses that β_γ is increasing, the second inequality just drops the constraint $h \geq f$, and the third inequality follows from $\mathcal{H}^m \subseteq \mathcal{H}$.

Fix $\varepsilon > 0$. By Condition (D) and Theorem 2.2 there exist $m_0 \in \mathbb{N}$ and $\gamma_0 > 0$ such that

$$\phi^m(f) \leq \phi(f) + \varepsilon \quad \text{and} \quad \phi(f) \leq \phi_{\theta, \gamma}(f) + \varepsilon$$

for all $m \geq m_0$ and $\gamma \geq \gamma_0$. This shows that

$$\phi(f) + \varepsilon + \frac{\beta(0)}{\gamma} \geq \phi^m(f) + \beta_\gamma(0) \geq \phi_{\theta, \gamma}^m(f) \geq \phi_{\theta, \gamma}(f) \geq \phi(f) - \varepsilon$$

for all $m \geq m_0$ and $\gamma \geq \gamma_0$, which shows that $\phi_{\theta, \gamma}^m(f) \rightarrow \phi(f)$ whenever $\min\{m, \gamma\} \rightarrow \infty$. \square

5.4 Proof of Lemma 3.3

- (a) From Hornik [34] it follows that \mathfrak{N}_{l_j, d_j} is dense in $C_b(\mathbb{R}^{d_j})$ with respect to $L^1(\mu)$ for every $\mu \in \mathcal{P}(\mathbb{R}^{d_j})$ and all $j = 1, \dots, J$. By the triangle inequality and boundedness of e_j , the first part of Condition (D) follows.

(b) If \mathcal{X} is compact, the condition is trivially satisfied. Hence assume that $\mathcal{X} = \mathbb{R}^d = \mathbb{R}^{d_1} \times \dots \times \mathbb{R}^{d_{J_0}}$ and $\pi_j = \text{pr}_j$, $e_j = 1$ for $j = 1, \dots, J_0 \leq J$, where pr_j is the projection from \mathbb{R}^d to the j -th marginal component in \mathbb{R}^{d_j} .

Let $\varepsilon > 0$ and $\mu \in \mathcal{P}(\mathcal{X})$. We first fix j , denote by $\mu^{(j)} := \mu \circ \text{pr}_j^{-1}$ and show that there exists a $h''_j \in \mathfrak{N}_{l_j, d_j}$ such that $1_{K_j^c} \leq h''_j$ and $\int_{\mathbb{R}^{d_j}} h''_j d\mu^{(j)} \leq 2\varepsilon$ for some compact subset K_j of \mathbb{R}^{d_j} . Without loss of generality, assume that $l_j = 1$. This can always be done since the function h''_j will be compact-valued and hence for multiple layers, the remaining layers beyond the first can simply approximate the identity function in the supremum norm.¹⁰ Fix $K_j = [-c, +c]^{d_j}$ such that $\mu^{(j)}(K_j^c) \leq \varepsilon/(4d_j)$. By assumption on φ , for each $i \in \{1, \dots, d_j\}$ there exist $\underline{a}_i, \underline{b}_i, \bar{a}_i, \bar{b}_i \in \mathbb{R}$ such that

$$\varphi(\underline{a}_i x_i + \underline{b}_i) + \varphi(\bar{a}_i x_i + \bar{b}_i) \begin{cases} \leq \varepsilon/(2d_j) & \text{for } x_i \in [-c, c] \\ \geq 1 - \varepsilon & \text{for } x_i \notin [-c - 1, c + 1]. \end{cases}$$

Then

$$h''_j := \sum_{i=1}^{d_j} \varphi(\underline{a}_i x_i + \underline{b}_i) + \varphi(\bar{a}_i x_i + \bar{b}_i) + \varepsilon \in \mathfrak{N}_{1, d_j, 2d_j} \subset \mathfrak{N}_{1, d_j}$$

satisfies $1_{\tilde{K}_j^c} \leq h''_j$ for the compact $\tilde{K}_j := [-c - 1, c + 1]^{d_j}$, as well as

$$\int_{\mathbb{R}^{d_j}} h''_j d\mu^{(j)} \leq \int_{\mathbb{R}^{d_j}} \frac{\varepsilon}{2} 1_{K_j} + 2d 1_{K_j^c} d\mu^{(j)} + \varepsilon \leq \frac{\varepsilon}{2} + 2d\mu^{(j)}(K_j^c) + \varepsilon \leq 2\varepsilon.$$

Now, define $h'' := \sum_{j=1}^d h''_j \circ \text{pr}_j \in \mathcal{H}^\infty$ and $K := \prod_{j=1}^{J_0} \tilde{K}_j \subset \mathcal{X}$, which is compact. Then one immediately gets $1_{K^c} \leq h''$ and $\int h'' d\mu \leq 2J_0\varepsilon$.

5.5 Proof of Proposition 3.7

(1) For one fix network $\mathfrak{N}_{l_j, d_j, m}$, the mapping $\xi \mapsto N_{l_j, d_j, m}(\xi)$ is pointwise continuous, i.e. it holds for $\xi_n \rightarrow \xi$ that $N_{l_j, d_j, m}(\xi_n)(x) \rightarrow N_{l_j, d_j, m}(\xi)(x)$ for all $x \in \mathbb{R}^{d_j}$, since φ is continuous. Further, since we assume that φ is bounded, the functions $N_{l_j, d_j, m}(\xi_n)$ are uniformly bounded and hence by dominated convergence one obtains $N_{l_j, d_j, m}(\xi_n) \rightarrow N_{l_j, d_j, m}(\xi)$ in $L^1(\mu)$ for all $\mu \in \mathcal{P}(\mathbb{R}^{d_j})$. By the triangle inequality, this continuity transfers to the mapping $(\xi_1, \dots, \xi_J) \mapsto \sum_{j=1}^J e_j N_{l_j, d_j, m}(\xi_j) \circ \pi_j$. Hence, we can write

$$\mathcal{H}^m = \{\eta(A) + a : A \in \mathcal{A}^m, a \in \mathbb{R}\}$$

¹⁰ More precisely, the function h''_j as given by one layer would be the input in the first component for the remaining layers, and the remaining layers approximate the continuous function $[-z, z]^m \ni x \mapsto x_1$ in the supremum norm, which is possible as shown in [34].

where $A \mapsto \eta(A)$ is continuous in $L^1(\mu_0)$ and $L^1(\theta)$, and \mathcal{A}^m is compact.
 (2) For every $\varepsilon > 0$ there exists $\eta(A) + a \in \mathcal{H}^m$ with $\eta(A) + a \geq f$ such that

$$\begin{aligned} \phi^m(f) + \varepsilon &\geq \int \eta(A) + a \, d\mu_0 \\ &= \lim_{\gamma \rightarrow \infty} \left\{ \int \eta(A) \, d\mu_0 + a + \int \beta_\gamma(f - \eta(A) - a) \, d\theta \right\} \\ &\geq \limsup_{\gamma \rightarrow \infty} \phi_{\theta, \gamma}^m(f) \end{aligned}$$

since $0 \leq \int \beta_\gamma(f - \eta(A) - a) \, d\theta \leq \beta_\gamma(0) = \frac{1}{\gamma} \beta(0)$.

On the other hand, let (γ_n) be a sequence in \mathbb{R}_+ with $\gamma_n \rightarrow \infty$. Our goal is to show that $\phi^m(f) \leq \liminf_{n \rightarrow \infty} \phi_{\theta, \gamma_n}^m(f)$. We assume that $\liminf_{n \rightarrow \infty} \phi_{\theta, \gamma_n}^m(f) < \infty$ otherwise there is nothing to prove. For every $n \in \mathbb{N}$ there exist $A^n \in \mathcal{A}^m$ and $a^n \in \mathbb{R}$ such that

$$\begin{aligned} \phi_{\theta, \gamma_n}^m(f) + \frac{1}{n} &\geq \int \eta(A^n) \, d\mu_0 + a^n + \int \beta_{\gamma_n}(f - \eta(A^n) - a^n) \, d\theta \\ &\geq \int \eta(A^n) \, d\mu_0 + \int f - \eta(A^n) + c \, d\theta \end{aligned} \tag{5.2}$$

since $\beta_{\gamma_n}(x) \geq x + c$ for all $n \in \mathbb{N}$ for some constant $c \in \mathbb{R}$. In particular, $\liminf_{n \rightarrow \infty} \phi_{\theta, \gamma_n}^m(f)$ is real-valued as $A \mapsto \int f - \eta(A) + c \, d\theta$ is a continuous function on the compact \mathcal{A}^m . By passing to a subsequence we may assume that $\lim_{n \rightarrow \infty} \phi_{\theta, \gamma_n}^m(f) = \liminf_{n \rightarrow \infty} \phi_{\theta, \gamma_n}^m(f)$, and $A^n \rightarrow A \in \mathcal{A}^m$ such that $\eta(A^n) \rightarrow \eta(A)$ in $L^1(\theta)$ and θ -a.s. as well as $\int \eta(A^n) \, d\mu_0 \rightarrow \int \eta(A) \, d\mu_0$. We next show that (a^n) is bounded. Suppose by way of contradiction that $a^n \rightarrow -\infty$. Since $\lim_{x \rightarrow \infty} \beta(x)/x = \infty$ and $f - \eta(A^n)$ is uniformly bounded by compactness of \mathcal{A}^m , it follows that

$$\int \eta(A^n) \, d\mu_0 + a^n + \beta_{\gamma_n}(f - \eta(A^n) - a^n) \rightarrow +\infty$$

Moreover, in view of (5.2) the sequence

$$\left(\int \eta(A^n) \, d\mu_0 + a^n + \beta_{\gamma_n}(f - \eta(A^n) - a^n) \right)^-$$

is uniformly integrable in $L^1(\theta)$. Hence, it follows from Fatou’s lemma that

$$\begin{aligned} +\infty &= \int \liminf_{n \rightarrow \infty} \left\{ \int \eta(A^n) \, d\mu_0 + a^n + \beta_{\gamma_n}(f - \eta(A^n) - a^n) \right\} \, d\theta \\ &\leq \liminf_{n \rightarrow \infty} \left\{ \int \eta(A^n) \, d\mu_0 + a^n + \int \beta_{\gamma_n}(f - \eta(A^n) - a^n) \, d\theta \right\} \\ &\leq \liminf_{n \rightarrow \infty} \phi_{\theta, \gamma_n}^m(f) < \infty \end{aligned}$$

which is the desired contradiction. This shows that (a^n) is bounded and by passing to a subsequence $a^n \rightarrow a \in \mathbb{R}$. Finally it follows from Fatou’s lemma that

$$\begin{aligned} \liminf_{n \rightarrow \infty} \phi_{\theta, \gamma_n}^m(f) &= \liminf_{n \rightarrow \infty} \left\{ \int \eta(A^n) d\mu_0 + a^n + \int \beta_{\gamma_n}(f - \eta(A^n) - a^n) d\theta \right\} \\ &\geq \int \eta(A) d\mu_0 + a + \int \beta_\infty(f - \eta(A) - a) d\theta \\ &= \int \eta(A) + a d\mu_0 \\ &\geq \phi^m(f) \end{aligned}$$

where $\beta_\infty(x) = 0$ if $x \leq 0$ and $\beta_\infty(x) = \infty$ if $x > 0$. The second inequality follows because $\eta(A) + a \geq f$ θ -a.s. as a consequence of the first inequality, $f, \eta(A) \in C_b(\mathcal{X})$ and θ is strictly positive. □

A Nonlinear version of the Daniell-Stone theorem

Let \mathcal{X} be a Polish space. Given a measurable function $\kappa : \mathcal{X} \rightarrow [1, \infty)$, we denote by $C_\kappa(\mathcal{X})$ the Stone vector lattice of all continuous functions $f : \mathcal{X} \rightarrow \mathbb{R}$ such that $|f|/\kappa$ is bounded. For instance, if κ is bounded one has $C_\kappa(\mathcal{X}) = C_b(\mathcal{X})$, or if $\kappa(x) = 1 + |x|$ on $\mathcal{X} = \mathbb{R}^d$ the space $C_\kappa(\mathbb{R}^d)$ contains all continuous functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ of linear growth. Further, let $ca_\kappa^+(\mathcal{X})$ be the set of all Borel measures μ on \mathcal{X} which satisfy $\int \kappa d\mu < \infty$. The following nonlinear version of the Daniell-Stone theorem follows directly from Proposition 1.1 in [15].

Proposition A.1 *Let $\phi : C_\kappa(\mathcal{X}) \rightarrow \mathbb{R}$ be an increasing¹¹ convex functional which is continuous from above, i.e. $\phi(f^n) \downarrow 0$ for every sequence (f^n) such that $f^n \downarrow 0$. Then, it has the dual representation*

$$\phi(f) = \max_{\mu \in ca_\kappa^+(\mathcal{X})} \left\{ \int f d\mu - \phi^*(\mu) \right\} \text{ for all } f \in C_\kappa(\mathcal{X}), \tag{A.1}$$

where the convex conjugate $\phi^* : ca_\kappa^+(\mathcal{X}) \rightarrow \mathbb{R} \cup \{+\infty\}$ is given by $\phi^*(\mu) = \sup_{f \in C_\kappa(\mathcal{X})} \{ \int f d\mu - \phi(f) \}$.

Continuity from above is strongly related to the concept of tightness, which in the context of risk measures was introduced by Föllmer and Schied, see [24]. Typical examples include transport type problems where tightness is imposed by marginal constraints, see e.g. Bartl et al. [5]. For extensions of the representation (A.1) to upper semicontinuous functions and related pricing-hedging dualities we refer to Cheridito et al. [16].

As an application we consider the superhedging functional

$$\phi(f) := \inf \left\{ \int h d\mu_0 : h \geq f \text{ for some } h \in \mathcal{H} \right\}$$

¹¹ $\phi(f) \geq \phi(g)$ whenever $f \geq g$.

on $C_\kappa(\mathcal{X})$, where $\mu_0 \in ca_\kappa^+(\mathcal{X})$ is a probability measure and $\mathcal{H} \subseteq C_\kappa(\mathcal{X})$ is a convex cone such that $\kappa \in \mathcal{H}$. Straightforward inspection shows that ϕ is a real-valued increasing convex functional on $C_\kappa(\mathcal{X})$. Further, if ϕ is continuous from above by Proposition A.1 it has the dual representation (A.1). Its convex conjugate is given by

$$\begin{aligned}\phi^*(\mu) &= \sup_{f \in C_\kappa(\mathcal{X})} \left\{ \int f d\mu - \inf_{\substack{h \in \mathcal{H}: \\ h \geq f}} \int h d\mu_0 \right\} \\ &= \sup_{h \in \mathcal{H}} \sup_{\substack{f \in C_\kappa(\mathcal{X}): \\ h \geq f}} \left\{ \int f d\mu - \int h d\mu_0 \right\} \\ &= \sup_{h \in \mathcal{H}} \left\{ \int h d\mu - \int h d\mu_0 \right\}.\end{aligned}\tag{A.2}$$

Since \mathcal{H} is a convex cone which contains the constants it follows that $\phi^*(\mu) = 0$ whenever $\mu \in ca_\kappa^+(\mathcal{X})$ is a probability measure such that $\int h d\mu = \int h d\mu_0$ for all $h \in \mathcal{H}$, and $\phi^*(\mu) = +\infty$ else. In particular, in case that $C_\kappa(\mathcal{X}) = C_b(\mathcal{X})$ we conclude the dual representation (2.2).

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: large-scale machine learning on heterogeneous systems (2015). Software available from www.tensorflow.org
- Alfonsi, A., Corbetta, J., Jourdain, B.: Sampling of probability measures in the convex order and approximation of martingale optimal transport problems (2017)
- Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein GAN. arXiv preprint [arXiv:1701.07875](https://arxiv.org/abs/1701.07875) (2017)
- Artzner, P., Delbaen, F., Eber, J.-M., Heath, D.: Coherent measures of risk. *Math. Financ.* **9**(3), 203–228 (1999)
- Bartl, D., Cheridito, P., Kupper, M., Tangpi, L.: Duality for increasing convex functionals with countably many marginal constraints. *Banach J. Math. Anal.* **11**(1), 72–89 (2017)
- Bartl, D., Kupper, M., Lux, T., Papapantoleon, A.: Sharpness of improved Fréchet-Hoeffding bounds: an optimal transport approach. arXiv preprint [arXiv:1709.00641](https://arxiv.org/abs/1709.00641) (2017)
- Beiglböck, M., Henry-Labordère, P., Penkner, F.: Model-independent bounds for option prices: a mass transport approach. *Financ. Stoch.* **17**(3), 477–501 (2013)
- Ben-Tal, A., Teboulle, M.: An old-new concept of convex risk measures: the optimized certainty equivalent. *Math. Financ.* **17**(3), 449–476 (2007)
- Benamou, J.-D., Carlier, G., Cuturi, M., Nenna, L., Peyré, G.: Iterative Bregman projections for regularized transportation problems. *SIAM J. Sci. Comput.* **37**(2), A1111–A1138 (2015)
- Bernard, C., Jiang, X., Wang, R.: Risk aggregation with dependence uncertainty. *Insurance* **54**, 93–108 (2014)
- Bernard, C., Rüschendorf, L., Vanduffel, S., Yao, J.: How robust is the value-at-risk of credit risk portfolios? *Eur. J. Financ.* **23**(6), 507–534 (2017)
- Bindini, U.: Smoothing operators in multi-marginal optimal transport. arXiv preprint [arXiv:1901.07407](https://arxiv.org/abs/1901.07407) (2019)
- Bühler, H., Gonon, L., Teichmann, J., Wood, B.: Deep hedging. arXiv preprint [arXiv:1802.03042](https://arxiv.org/abs/1802.03042) (2018)

14. Carlier, G., Duval, V., Peyré, G., Schmitzer, B.: Convergence of entropic schemes for optimal transport and gradient flows. *SIAM J. Math. Anal.* **49**(2), 1385–1418 (2017)
15. Cheridito, P., Kupper, M., Tangpi, L.: Representation of increasing convex functionals with countably additive measures. arXiv preprint [arXiv:1502.05763](https://arxiv.org/abs/1502.05763) (2015)
16. Cheridito, P., Kupper, M., Tangpi, L.: Duality formulas for robust pricing and hedging in discrete time. *SIAM J. Financ. Math.* **8**(1), 738–765 (2017)
17. Cominetti, R., San Martín, J.: Asymptotic analysis of the exponential penalty trajectory in linear programming. *Math. Program.* **67**(1–3), 169–187 (1994)
18. Cuturi, M.: Sinkhorn distances: lightspeed computation of optimal transport. In: *Advances in neural information processing systems*, pp. 2292–2300 (2013)
19. Eckstein, S., Kupper, M., Pohl, M.: Robust risk aggregation with neural networks. arXiv preprint [arXiv:1811.00304](https://arxiv.org/abs/1811.00304) (2018)
20. Ekren, I., Soner, H.M.: Constrained optimal transport. *Arch. Ration. Mech. Anal.* **227**(3), 929–965 (2018)
21. Embrechts, P., Puccetti, G., Rüschendorf, L.: Model uncertainty and VaR aggregation. *J. Bank. Financ.* **37**(8), 2750–2764 (2013)
22. Feizi, S., Suh, C., Xia, F., Tse, D.: Understanding GANs: the LQG setting. arXiv preprint [arXiv:1710.10793](https://arxiv.org/abs/1710.10793) (2017)
23. Flury, B.D.: Acceptance-rejection sampling made easy. *SIAM Rev.* **32**(3), 474–476 (1990)
24. Föllmer, H., Schied, A.: *Stochastic Finance: An Introduction in Discrete Time*. Walter de Gruyter, Berlin (2011)
25. Galichon, A., Henry-Labordere, P., Touzi, N.: A stochastic control approach to no-arbitrage bounds given marginals, with an application to lookback options. *Ann. Appl. Probab.* **24**(1), 312–336 (2014)
26. Genevay, A., Cuturi, M., Peyré, G., Bach, F.: Stochastic optimization for large-scale optimal transport. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 29, pp. 3440–3448. Curran Associates, Inc, New York (2016)
27. Genevay, A., Cuturi, M., Peyré, G., Bach, F.: Stochastic optimization for large-scale optimal transport. In: *Advances in Neural Information Processing Systems*, pp. 3440–3448 (2016)
28. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in neural information processing systems*, pp. 2672–2680 (2014)
29. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.: Improved training of Wasserstein GANs. arXiv preprint [arXiv:1704.00028](https://arxiv.org/abs/1704.00028) (2017)
30. Guo, G., Obloj, J.: Computational methods for martingale optimal transport problems. arXiv preprint [arXiv:1710.07911](https://arxiv.org/abs/1710.07911) (2017)
31. Gurobi, L.: *Optimization: Gurobi optimizer reference manual* (2018)
32. Henry-Labordere, P.: Automated option pricing: numerical methods. *Int. J. Theor. Appl. Financ.* **16**(08), 1350042 (2013)
33. Horger, F., Würfl, T., Christlein, V., Maier, A.: Deep learning for sampling from arbitrary probability distributions. arXiv preprint [arXiv:1801.04211](https://arxiv.org/abs/1801.04211) (2018)
34. Hornik, K.: Approximation capabilities of multilayer feedforward networks. *Neural Netw.* **4**(2), 251–257 (1991)
35. Kantorovich, L.V.: On the translocation of masses. *Dokl. Akad. Nauk SSSR* **37**, 199–201 (1942)
36. Kellerer, H.G.: Duality theorems for marginal problems. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* **67**(4), 399–432 (1984)
37. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
38. Léonard, C.: A survey of the Schrödinger problem and some of its connections with optimal transport. *Discret. Continuous Dyn. Syst. A* **34**(4), 1533–1574 (2014)
39. Lux, T., Papapanoleon, A.: Improved Fréchet-Hoeffding bounds on d -copulas and applications in model-free finance. *Ann. Appl. Probab.* **27**(6), 3633–3671 (2017)
40. Peng, S.: G-expectation, G-Brownian motion and related stochastic calculus of Itô type. In: *Stochastic analysis and applications*, pp. 541–567. Springer, New York (2007)
41. Pflug, G.C., Pichler, A.: *Multistage stochastic optimization*. Springer, New York (2014)
42. Pflug, G.C., Pohl, M.: A review on ambiguity in stochastic portfolio optimization. *Set-Valued Var. Anal.* **23**(1), 11–25 (2017)

43. Puccetti, G., Rüschendorf, L.: Computation of sharp bounds on the distribution of a function of dependent risks. *J. Comput. Appl. Math.* **236**(7), 1833–1840 (2012)
44. Schmitzer, B.: Stabilized sparse scaling algorithms for entropy regularized transport problems. arXiv preprint [arXiv:1610.06519](https://arxiv.org/abs/1610.06519) (2016)
45. Seguy, V., Damodaran, B.B., Flamary, R., Courty, N., Rolet, A., Blondel, M.: Large-scale optimal transport and mapping estimation. arXiv preprint [arXiv:1711.02283](https://arxiv.org/abs/1711.02283) (2017)
46. Smith, S.L., Kindermans, P.-J., Le, Q.V.: Don't decay the learning rate, increase the batch size. arXiv preprint [arXiv:1711.00489](https://arxiv.org/abs/1711.00489) (2017)
47. Solomon, J., De Goes, F., Peyré, G., Cuturi, M., Butscher, A., Nguyen, A., Du, T., Guibas, L.: Convolutional Wasserstein distances: efficient optimal transportation on geometric domains. *ACM Trans. Gr. (TOG)* **34**(4), 66 (2015)
48. Strassen, V.: The existence of probability measures with given marginals. *Ann. Math. Stat.* **36**(2), 423–439 (1965)
49. Vallender, S.: Calculation of the Wasserstein distance between probability distributions on the line. *Theory Probab. Appl.* **18**(4), 784–786 (1974)
50. Villani, C.: *Optimal Transport: Old and New*, vol. 338. Springer, New York (2008)
51. Weinan, E., Han, J., Jentzen, A.: Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Commun. Math. Stat.* **5**(4), 349–380 (2017)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.